

MuMAX: a new high-performance micromagnetic simulation tool

A. Vansteenkiste^a, B. Van de Wiele^b

^aDepartment of Solid State Sciences, Ghent University, Krijgslaan 281-S1, B9000 Gent, Belgium.

^bDepartment of Electrical Energy, Systems and Automation, Ghent University, Sint Pietersnieuwstraat 41, B-9000 Ghent, Belgium.

Abstract

We present MuMAX, a general-purpose micromagnetic simulation tool running on Graphical Processing Units (GPUs). MuMAX is designed for high performance computations and specifically targets large simulations. In that case speedups of over a factor 100× can easily be obtained compared to the CPU-based OOMMF program developed at NIST. MuMAX aims to be general and broadly applicable. It solves the classical Landau-Lifshitz equation taking into account the magnetostatic, exchange and anisotropy interactions, thermal effects and spin-transfer torque. Periodic boundary conditions can optionally be imposed. A spatial discretization using finite differences in 2 or 3 dimensions can be employed. MuMAX is publicly available as open source software. It can thus be freely used and extended by community. Due to its high computational performance, MuMAX should open up the possibility of running extensive simulations that would be nearly inaccessible with typical CPU-based simulators.

Keywords: micromagnetism, simulation, GPU

PACS: 75.78.Cd, 02.70.Bf

1. Introduction

Micromagnetic simulations are indispensable tools in the field of magnetism research. Hence, micromagnetic simulators like, e.g., OOMMF [1], magpar [2] and Nmag [3] are widely used. These tools solve the Landau-Lifshitz equation on regular CPU hardware. Due to the required fine spatial and temporal discretizations, such simulations can be very time consuming. Limited computational resources therefore often limit the full capabilities of the otherwise successful micromagnetic approach.

There is currently a growing interest in running numerical calculations on graphical processing units (GPUs) instead of CPUs. Although originally intended for purely graphical purposes, GPUs turn out to be well-suited for high-performance, general-purpose calculations. Even relatively cheap GPUs can perform an enormous amount of calculations in parallel. E.g., the nVIDIA GTX580 GPU used for this work costs less than \$500 and delivers 1.5 trillion floating-point operations (Flops) per second, about 2 orders of magnitude more than a typical CPU.

However, in order to employ this huge numerical power programs need to be written specifically for GPU hardware, using the programming languages and tools provided by the GPU manufacturer, and the code also needs to handle many hardware-specific technicalities. Additionally, the used algorithms need to be expressed in a highly parallel manner, which is not always easily possible.

Other groups have already implemented micromagnetic simulations on GPU hardware and report considerable speedups

compared to a CPU-only implementation [4, 5, 6]. At the time of writing, however, none of these implementations is freely available. MuMAX, on the other hand, is available as open source software and can be readily used by anyone. Its performance also compares favorably to these other implementations.

2. Methods

Since the micromagnetic theory describes the magnetization as a continuum field $\mathbf{M}(\mathbf{r}, t)$, the considered magnetic sample is discretized in cuboidal finite difference (FD) cells with a uniform magnetization. The time evolution of the magnetization in each cell is given by the Landau-Lifshitz equation

$$\frac{\partial \mathbf{M}(\mathbf{r}, t)}{\partial t} = -\frac{\gamma}{1 + \alpha^2} \mathbf{M}(\mathbf{r}, t) \times \mathbf{H}_{eff}(\mathbf{r}, t) - \frac{\alpha\gamma}{M_s(1 + \alpha^2)} \mathbf{M}(\mathbf{r}, t) \times (\mathbf{M}(\mathbf{r}, t) \times \mathbf{H}_{eff}(\mathbf{r}, t)). \quad (1)$$

Here, M_s is the saturation magnetization, γ the gyromagnetic ratio and α the damping parameter. The continuum effective field \mathbf{H}_{eff} has several contributions that depend on the magnetization, the externally applied field and the material parameters of the considered sample. When timestepping equation (1) the effective field is evaluated several times per time step. Hence, the efficiency of micromagnetic software depends on the efficient evaluation of the different effective field terms at the one hand and the application of efficient time stepping schemes on the other hand. MuMAX combines both with the huge computational power of GPU hardware.

Email address: Arne.Vansteenkiste@ugent.be (A. Vansteenkiste)

2.1. Effective field terms

In the present version of MuMAX, the effective field can have 5 different contributions: the magnetostatic field, the exchange field, the applied field, the anisotropy field and the thermal field. In what follows we present these terms and comment on their optimized implementation.

2.1.1. Magnetostatic field

The magnetostatic field \mathbf{H}_{ms} accounts for the long-range interaction throughout the complete sample

$$\mathbf{H}_{ms}(\mathbf{r}) = -\frac{1}{4\pi} \int_V \nabla \nabla \frac{1}{|\mathbf{r} - \mathbf{r}'|} \cdot \mathbf{M}(\mathbf{r}') d\mathbf{r}'. \quad (2)$$

Since the magnetostatic field in one FD cell depends on the magnetization in all other FD cells, the calculation of \mathbf{H}_{ms} is the most time-consuming part of a micromagnetic simulation. The chosen method for this calculation is thus decisive for the performance of the simulator. Therefore, we opted for a fast Fourier transform (FFT) based method. In this case, the convolution structure of (2) is exploited. By applying the convolution theorem, the convolution is accelerated by first Fourier transforming the magnetization, then multiplying this result with the Fourier-transform of the convolution kernel and finally inverse transforming this product to obtain the magnetostatic field. The overall complexity of this method is $O(N \log N)$, as it is dominated by the FFTs.

Methods with even lower complexity exist as well. The fast multipole method, e.g., only has complexity $O(N)$, but with such a large pre-factor that in most cases the FFT method remains much faster [7].

A consequence of the FFT method is that the magnetic moments must lie on a regular grid. This means that a finite difference (FD) spatial discretization has to be used: space is divided into equal cuboid cells. This method is thus most suited for rectangular geometries. Other shapes have to be approximated in a staircase-like fashion. However, thanks to the speedup offered by MuMAX's, smaller cells may be used to improve this without excessive performance penalties.

The possibility of adding periodic boundary conditions in one or more directions is also included in the software. This is done by adding a sufficiently large number of periodic images to the convolution kernel. The application of periodic boundary conditions has a positive influence on the computational time since the magnetization data does not need to be zero padded in the periodic directions, which roughly halves the time spend on FFTs for every periodic direction.

2.1.2. Exchange field

The exchange interaction contributes to the effective field in the form of a laplacian of the magnetization

$$\mathbf{H}_{exch} = \frac{2A}{\mu_0 M_s} \nabla^2 \mathbf{m}, \quad (3)$$

with A the exchange stiffness. In discretized form, this can be expressed as a linear combination of the magnetization of a cell and a number of its neighbors. MuMAX uses a 6-neighbor

scheme, similar to [8]. In the case of 2D simulations (only one FD cell in the z-direction), this method automatically reduces to a 4-neighbor scheme.

The exchange field calculation is included in the magnetostatic field routines by simply adding the kernel describing the exchange interaction to the magnetostatic kernel. In this way, the exchange calculation is essentially free, as only one joint convolution product is needed to simultaneously evaluate both the magnetostatic and exchange fields. Moreover, by introducing the exchange contribution in the magnetostatic field kernel periodic boundary conditions are directly accounted for if applicable.

2.1.3. Other effective field terms

Next to the above mentioned interaction terms and the applied field contribution, MuMAX provides the ability to include magnetocrystalline anisotropy. Currently, uniaxial and cubic anisotropy are available. The considered anisotropy energies are

$$\phi_{ani} = K_u \sin^2 \theta \quad (4)$$

and

$$\begin{aligned} \phi_{ani}(\mathbf{r}) = & K_1 [\alpha_1^2(\mathbf{r})\alpha_2^2(\mathbf{r}) + \alpha_2^2(\mathbf{r})\alpha_3^2(\mathbf{r}) + \alpha_1^2(\mathbf{r})\alpha_3^2(\mathbf{r})] \\ & + K_2 [\alpha_1^2(\mathbf{r})\alpha_2^2(\mathbf{r})\alpha_3^2(\mathbf{r})] \end{aligned} \quad (5)$$

for uniaxial and cubical anisotropy respectively. Here, K_u and (K_1, K_2) are the uniaxial and cubical anisotropy constants, θ is the angle between the local magnetization and uniaxial anisotropy axis and α_i ($i = 1, 2, 3$) are the direction cosines between the local magnetization and the cubic easy magnetization axes.

Furthermore, thermal effects are included by means of a fluctuating thermal field

$$\mathbf{H}_{th} = \boldsymbol{\eta}(\mathbf{r}, t) \sqrt{\frac{2\alpha k_B T}{\gamma \mu_0 M_s V \delta t}} \quad (6)$$

which is added to the effective field \mathbf{H}_{eff} according to [9]. In (6), k_B is the Boltzmann constant, V is the volume of a FD cell, δt is the used time step and $\boldsymbol{\eta}(\mathbf{r}, t)$ is a stochastic vector whose components are Gaussian random numbers, uncorrelated in space and time with zero mean value and dispersion 1.

2.1.4. Spin-transfer torque

The spin-transfer torque interaction describes the influence of electrical currents on the local magnetization. Possible applications are spin-transfer torque random access memory [10] and racetrack memory [11]. MuMAX incorporates the spin-transfer torque description developed by Berger [12], refined

by Zhang and Li [13]

$$\begin{aligned} \frac{\partial \mathbf{M}}{\partial t} = & -\frac{\gamma}{1+\alpha^2} \mathbf{M} \times \mathbf{H}_{eff} \\ & -\frac{\alpha\gamma}{M_s(1+\alpha^2)} \mathbf{M} \times (\mathbf{M} \times \mathbf{H}_{eff}) \\ & -\frac{b_j}{M_s^2(1+\alpha^2)} \mathbf{M} \times (\mathbf{M} \times (\mathbf{j} \cdot \nabla) \mathbf{M}) \\ & -\frac{b_j}{M_s(1+\alpha^2)} (\xi - \alpha) \mathbf{M} \times (\mathbf{j} \cdot \nabla) \mathbf{M}. \end{aligned} \quad (7)$$

Here, ξ is the degree of non-adiabicity and b_j is the coupling constant between the current density \mathbf{j} and the magnetization

$$b_j = \frac{P\mu_B}{eM_s(1+\xi^2)}, \quad (8)$$

with P the polarization of the current density, μ_B the Bohr magneton and e the electron charge.

2.2. Time integration schemes

MuMAX provides a range of Runge-Kutta (RK) methods to integrate the Landau-Lifshitz equation. Currently the user can select between the following options:

- RK1: Euler's method
- RK2: Heun's method
- RK12: Heun-Euler (adaptive step)
- RK3: Kutta's method
- RK23: Bogacki-Shampine (adaptive step)
- RK4: Classical Runge-Kutta method
- RKCK: Cash-Karp (adaptive step)
- RKDP: Dormand-Prince (adaptive step)

The adaptive step methods adjust the time step based on a maximum tolerable error per integration step that can be set by the user. The other methods can use either a fixed time step or a fixed maximum variation of \mathbf{m} per step. Depending on the needs of the simulation, a very accurate but relatively slow high-order solver (e.g. RKDP) or a less accurate but fast solver (e.g. RK12) can be chosen. Additionally, MuMAX incorporates the semi-analytical methods described in [14]. These methods are specifically tailored to the Landau-Lifshitz equation.

3. GPU-optimized implementation

Since various CPU based micromagnetic tools—well suited for relatively small micromagnetic problems—are already available, we mainly concentrated on optimizing MuMAX for running very large simulations on GPUs. Nevertheless, the code can also run in CPU-mode, with multi-threading modalities enabled. In this way one can get familiar with the capabilities of MuMAX before a high-end GPU has to be purchased.

The GPU-specific parts of MuMAX have been developed using NVIDIA's CUDA platform. The low-level, performance-critical functions that have to interact directly with the GPU are written in C/C++. Counterparts of these functions for the CPU are implemented as well and use FFTW [15] and multi threading. The high-level parts of MuMAX are implemented in "safe" languages including Java, Go and Python. This part is independent of the underlying GPU/CPU hardware. In what follows we will only elaborate on the GPU-optimized implementation of the low-level functions.

3.1. General precautions

A high-end GPU has its own dedicated memory with a high bandwidth (typically a few hundred GB/s) which enables fast reads and writes on the GPU itself. Communication with the CPU on the other hand is much slower since this takes place over a PCI express bus with a much lower bandwidth (typically a few GB/s). Therefore, our implementation keeps as much data as possible in the dedicated GPU memory, avoiding CPU-GPU communication. The only large data transfers occur at initialization time and when output is saved to disk. The CPU thus only instructs the GPU which subroutines to launch. Hence, the GPU handles all the major computational jobs.

On the GPU, an enormous number of threads can run in parallel, each performing a small part of the computations. E.g., the GTX580 GPU used for this work has 512 computational cores grouped in 16 multiprocessors, resulting in total number of 16384 available parallel threads. However, this huge parallel power is only optimally exploited when the code is adapted to the specific GPU architecture. E.g.: threads on the same multiprocessor ("thread warps") should only access the GPU memory in a coalesced way and should ideally perform the same instructions. When coalesced memory access is not possible, the so-called *shared memory* should be used instead of the global GPU memory. This memory is faster and has better random-access properties but is very scarce. Our implementation takes into account all these technicalities, resulting in a very high performance.

3.2. GPU-optimization of the convolution product

Generally, most computational time goes to the evaluation of the convolution product defined by the magnetostatic field. When using fast Fourier transforms (FFTs), the computations enhance three different stages: (i) forward Fourier transforming the magnetization data that is zero padded in the non-periodic directions, (ii) point-by-point multiplying the obtained data with the Fourier transformed magnetostatic field kernel, (iii) inverse Fourier transforming the resulting magnetostatic field data. The careful implementation of these three stages determines the efficiency of the convolution product and, more general, of the micromagnetic code.

In the first place, the efficiency of this convolution process is safeguarded by ensuring that the matrices defined by the magnetostatic field kernel are completely symmetrical. Consequently, the Fourier transformed kernel data is purely real. The

absence of the imaginary part leads to smaller memory requirements as well as a much faster evaluation of the point-by-point multiplications – step (ii).

Furthermore, our GPU implementation of the fast Fourier transforms, which internally uses the CUDA "CUFFT" library, is specifically optimized for micromagnetic applications. The general 3D real-to-complex Fourier transform (and its inverse) available in the CUFFT library is replaced by a more efficient implementation in which the set of 1D transforms in the different directions are performed separately. This way, Fourier transforms on arrays containing only zeros resulting from the zero padding are avoided. In each dimension, the set of 1D Fourier transforms are performed on contiguous data points resulting in the coalesced reading and writing of the data. As a drawback, the transposition of the data between a set of Fourier transforms in one and another dimension is needed.

In a straight forward implementation of the required matrix transposes, the read and write instructions can not be both performed in a coalesced way since either the input data or the transposed data is not contiguous in global memory. Therefore, the input data is divided in blocks and copied to shared memory assigned to a predefined number of GPU threads. There, the data block is transposed and copied in a coalesced way back to the global GPU memory space. By inventively using the large number of zero arrays –in the non-periodic case– this transpose process can be done without (for 2D) or with only limited (3D) extra memory requirements. The different sets of Fourier transforms in this approach are performed using the 1D FFT routines available in the CUFFT library. This implementation outperforms the general 3D real-to-complex Fourier transform available in the CUFFT library while the built-in 2D real-to-complex is only faster for small dimensions (for square geometries: smaller than 512x512 FD cells). This approach ensures the efficient evaluation of steps (i) and (iii) of the convolution.

3.3. Floating point precision

GPUs are in general better suited for single-precision than double-precision arithmetic. Double-precision performance is not only much slower due to the smaller number of arithmetic units, but also requires twice the amount of memory. Since GPU's typically have limited memory and FFT methods are relatively memory-intensive, we opted to use single-precision exclusively.

While, e.g., the finite element method used by Kakay et al. relies heavily on double precision to obtain an accurate solution [4], our implementation is designed to remain accurate even at single precision. First, all quantities are internally stored in units that are well adapted to the problem. More specifically, we choose units so that $\mu_0 = \gamma_0 = M_s = A = 1$. This avoids that any other quantity in the simulation becomes exceptionally large or small—which could cause a loss of precision due to saturation errors. The conversion to and from internal units is performed transparently to the user. Secondly, we avoid numerically instable operations like, e.g., subtracting nearly equal numbers. This avoids that small rounding of errors get amplified. Finally, and most importantly, we restrict the size of the FFTs to numbers where the CUFFT implementation is most

accurate: $2^n \times \{1, 3, 5 \text{ or } 7\}$. Hence sometimes a slightly larger number of FD cells than strictly necessary is used to meet this requirement. Fortunately this has no adverse effect on the performance since CUFFT FFTs with these sizes also happen to be exceptionally fast (see below). In this way, the combined error introduced by the forward+inverse FFT was found to be only of the order of $O(10^{-6})$, as opposed to a typical error of $O(10^{-4})$ for other FFT sizes (estimated from the error on transforming random data back and forth). Thanks to these precautions we believe that our implementation should be sufficiently accurate for most practical applications. Indeed, the uncertainty on material parameters alone is usually much larger than the FFT error of 10^{-6} .

4. Validation

In order to validate our software, we tested the reliability of the code by simulating several standard problems. These standard problems are constructed such that all different contributions in the considered test case influence the magnetization processes significantly. A correct simulation of standard problems can be considered as the best possible indication of the validity of the developed software. In what follows, we consider standard problems constructed for testing static simulations, dynamic simulations and dynamic simulations incorporating spin-transfer torque.

4.1. static standard problem

The μ MAG standard problem #2 [16] aims at testing quasi static simulations. A cuboid with dimensions $5d \times d \times 0.1d$ is considered. Since only magnetostatic and exchange interactions are included, the resulting static properties only depend on the scaled parameter d/l_{ex} , with l_{ex} the exchange length. The starting configuration is saturation along the $[1, 1, 1]$ axis, which is relaxed to the remanent state. This was done by solving the Landau-Lifshitz equation with a high damping parameter $\alpha = 1$.

The number of FD cells was chosen depending on the size of nanostructure, making sure the cell size remained below the exchange length. For $d/l_{ex} \leq 10$, single-domain states with nearly full saturation along the long axis were found, while for large geometries an S-state occurred. The MuMAX simulations considered 200 values for d/l_{ex} . On the GPU a total simulation time of 3'21" was needed to complete the 200 individual simulations, compared to 34'30" on the CPU. The GPU speedup is here limited by the relatively small simulation sizes (cfr. Fig. 7).

Figure 1 shows the remanent magnetization in function of the ratio d/l_{ex} . The values obtained with MuMAX coincide well with those of other authors [17, 18], validating MuMAX for static micromagnetic problems.

4.2. dynamic standard problem

The μ MAG standard problem #4 [16] aims at testing the description of the dynamic magnetization processes by considering the magnetization reversal in a thin film with dimensions




Figure 1: Standard problem #2. Remanent magnetization along the x -axis (left axis) and along the y -axis (right axis) as a function of the scaling parameter d . The full line represents the simulation results from MuMax, while the circles represent simulation points obtained from McMichael et al. [17] and from Donahue et al. [18]

500 nm \times 125 nm \times 3 nm. Starting from an initial equilibrium S-state, two different fields are applied. In this problem only the exchange and magnetostatic interactions are considered (exchange stiffness $A = 1.3 \times 10^{-11}$ J/m, saturation magnetization $M_s = 8.0 \times 10^5$ Am $^{-1}$). When relaxing to the initial equilibrium S-state, a damping constant equal to 1 is used while during the reversal itself a damping constant of 0.02 is applied, according to the problem definition. As proposed in the standard problem, we show in Figs. 2 and 3 the evolution of the average magnetization components together with the reference magnetization configuration at the time point when $\langle M_x \rangle$ crosses zero for the first time, for field 1 ($\mu_0 H_x = -24.6$ mT, $\mu_0 H_y = 4.3$ mT, $\mu_0 H_z = 0.0$ mT) and field 2 ($\mu_0 H_x = -35.5$ mT, $\mu_0 H_y = -6.3$ mT, $\mu_0 H_z = 0.0$ mT). A discretization using $128 \times 32 \times 1$ FD cells and the RK23 time stepping scheme with a time step around 600 fs (dynamically adapted during the simulation) was used. The relatively large time steps used to solve this standard problem demonstrate that MuMax incorporates robust time stepping schemes with accurate adaptive step mechanisms. The adaptive step algorithms ensure optimal time step lengths and thus reduce the number of field evaluations, speeding up the simulation. Here, a total time of only 2.5 seconds was needed to finish this simulation on the GPU compared to 16 seconds on the CPU. In this case the speedup on GPU is limited due to the small number of FD cells. When the simulation is repeated with a finer discretization of $256 \times 64 \times 2$ cells, on the other hand, the GPU speedup already becomes more pronounced: the simulation finishes in 46 seconds on the GPU but takes 20'32" on the CPU.

From Figs. 2 and 3 it is clear that the results obtained with MuMax are well within the spread of the curves obtained by other authors. Also the magnetization plots are in close agreement with those available at the μ Mag website [16].





Figure 2: (top) Time evolution of the average magnetization during the reversal considered in μ Mag standard problem #4, field1. The results obtained with MuMax (black) lie well within the spread of the reference solutions (grey), taken from [16]. (bottom) Magnetization configuration when $\langle M_x \rangle$ crosses the zero magnetization for the first time.





Figure 3: (top) Time evolution of the average magnetization during the reversal considered in μ Mag standard problem #4, field2. This field was chosen to cause a bifurcation point to make the different solutions diverge. (bottom) Magnetization configuration when $\langle M_x \rangle$ crosses the zero magnetization point for the first time.




Figure 4: Time evolution of the average in plane magnetization during the first 8 ns of the spin-transfer torque standard problem. To facilitate the visual comparison of our results with [19], the average magnetization is expressed in [A/m] and the same axes ratios are chosen.

4.3. Spin-transfer torque standard problem [19]

μMag does not propose any standard problems that include spin-transfer torque. Therefore we rely on a standard problem proposed by M. Najafi et al. [19] to check the validity of the spin-transfer torque description implemented in MuMax. The standard problem considers a permalloy sample ($A = 1.3 \times 10^{-11} \text{ J/m}$, $M_s = 8.0 \times 10^5 \text{ Am}^{-1}$) with dimensions $100 \text{ nm} \times 100 \text{ nm} \times 10 \text{ nm}$. The initial equilibrium magnetization state is a predefined vortex, positioned in the center of the sample and relaxed without any spin-transfer torque interaction ($\alpha = 1.0$). Once relaxed, an homogeneous spin-polarized dc current $\mathbf{j} = 10^{12} \text{ Am}^{-2}$ along the x -axis is applied on the sample. Now, α is 0.1 and the degree of non-adiabaticity ξ is 0.05, see expression (7). Under these circumstances, the vortex center moves towards a new equilibrium position. The time evolution of the average in plane magnetization and the magnetization configuration at $t=0.73 \text{ ns}$ are shown respectively in Fig. 4 and Fig. 5. The results are in good agreement with those presented in reference [19]. With a discretization of 128×128 FD cells, 10 minutes of simulation time were needed to obtain the presented data.

5. Performance

The performance of MuMax on the CPU is roughly comparable to OOMMF. The CPU performance is thus good, but our main focus is optimizing the GPU code. Special attention went to fully exploiting the numerical power of the GPU while focussing on the time- and memory-efficient simulation of large micromagnetic problems. Figure 6 shows the time required to take one time step with the Euler method (i.e. effective field evaluation, evaluation of the LL-equation and magnetization update) on CPU (1 core) and on GPU for 2D and 3D simulations.

In both the 2D and 3D case, speedups of up two orders of magnitude are obtained for large dimensions. For smaller geometries, the speedups decrease but remain significant. This




Figure 5: Magnetization configuration at $t=0.73 \text{ ns}$ as found during the simulation of a standard problem incorporating spin-transfer torque [19]. The vortex core evolves towards a new equilibrium state under influence of a spin-polarized dc current along the horizontal direction. This figure is rendered with the built-in graphics features present in MuMax.

can be understood by the fact that in these simulations not enough FD cells are considered to have all $O(10^4)$ available threads at work at the same time. Hence, the computational power is not fully exploited. Furthermore, Fig. 6 shows that the CPU performance as well as the GPU performance does not follow a smooth curve. This is a consequence of the FFTs which are most efficient for powers of two, possibly multiplied with one small prime (in the benchmarks shown in Fig. 6, the default rounding to these optimal sizes is not performed). In the GPU implementation this is even more the case than in the CPU implementation. E.g., the 2D simulation with dimensions 992×992 is five times slower than the 2D simulation with dimensions 1024×1024 . This shows that not only for accuracy reasons, but also for time efficiency reasons, it is most advantageous to restrict the simulations domain to the optimal dimensions defined by $2^n \times \{1, 3, 5 \text{ or } 7\}$. Because of the extreme impact on the performance of MuMax, we opted to standardly rescale the size of the FD cells such that the dimensions are rounded of to one of these optimal sizes.

Due to the typical architecture of GPUs and the nature of the FFT algorithm, simulations of geometries with power of two sizes run extremely fast on GPU. Table 1 gives an overview of the speedups for these sizes for the 2D and 3D case and Fig. 7 shows the speedup obtained for these power of two sizes by MuMax compared to the OOMMF code. Both comparisons result in speedups larger than a factor 100. This means that simulations that used to take several hours can now be performed in minutes. The comparison between the speedups shown in Table 1 and Fig. 7 further show that our CPU implementation has indeed a comparable efficiency regarding to OOMMF. The immense speedups evidence the fact that MuMax can indeed open completely new research opportunities in micromagnetic modelling.




Figure 6: Time required to perform one time step using the Euler method for (top) 2D geometries with varying dimensions $N \times N$ and (bottom) 3D geometries with varying dimensions $N \times N \times N$. The CPU computations are performed on a 2.8 GHz intel core i7-930 processor, while the GPU computations are performed on nVIDIA GTX580 GPU hardware.

Table 1: Time needed to take one time step with the Euler method on CPU and GPU for 2D geometries (top) and 3D geometries (bottom).

size	CPU time (ms)	GPU time (ms)	speedup
32^2	0.633	0.59	$\times 1.07$
64^2	3.092	0.60	$\times 5.1$
128^2	6.739	0.69	$\times 9.7$
256^2	59.90	1.11	$\times 18$
512^2	266.8	2.75	$\times 47$
1024^2	1166	9.07	$\times 128$
2048^2	5233	35.78	$\times 146$
8^3	0.8492	0.79	$\times 1.07$
16^3	4.066	1.03	$\times 3.9$
32^3	36.14	1.70	$\times 21$
64^3	489.6	5.52	$\times 88$
128^3	4487	35.42	$\times 126$




Figure 7: Speedup obtained with MuMax running on a GTX580 GPU compared to OOMMF on a 2.8GHz core i7-930 CPU. The 2D and 3D geometries have sizes $N \times N$ and $N \times N \times N$ respectively. The lowest speedup for the $16 \times 16 \times 16$ case – an unusually small simulation – is still a factor 4.

6. How to use MuMax

MuMAX is released as open source software under the GNU General Public License (GPL) v.3 and can thus be freely used by the community. In addition to the terms of the GPL, we kindly ask to acknowledge the authors in any publication or derivative software that uses MuMAX, by citing this paper. The MuMAX source code can be obtained via <http://dynamat.ugent.be/mumax>. To use the software, a PC with a "CUDA capable" nVIDIA GPU and a recent 64-bit Linux installation is required.

A MuMAX simulation is entirely specified by an input file passed via the command-line. I.e., once the input file is written, no further user interaction is necessary to complete the simulation. This allows, for instance, to run large batches of simulations unattended. Nevertheless, the progress of a simulation can easily be checked: the number of time steps taken, total simulated time, etc. is reported in the terminal, PNG images of the magnetization state can be output on-the-fly, graphs of the average magnetization can easily be obtained while the simulation is running, etc. Furthermore, MuMAX's output format is compatible with OOMMF, enabling the use of existing post-processing tools to visualize and analyze the output. Built-in tools for output processing are available as well. The 3D vector field in Fig. 5, e.g., is rendered by MuMAX's tools.

MuMAX input files can be written in Python. This offers powerful control over the simulation flow and output. As an example, the code snippet below simulates an MRAM element as in standard problem 2 (see section 4.1). Starting from a uniform state in the $+x$ direction, it scans the field B in small steps until the point of coercivity. This illustrates how easily complex simulations can be defined.

```

from mumax import *

msat(800e3)
aexch(1.3e-11)
partsize(500e-9, 50e-9, 5e-9)
uniform(1, 0, 0)

B = 0
while avg_m('x') > 0:
    staticfield(-B, 0, 0)
    relax()
    B += 1e-4
# B now holds the coercitive field
save('m', 'binary')

```

Figure 8: MuMax input file snippet illustrating a simulation specification in Python. After initialization, a field in the $-x$ direction is stepped until the average magnetization along x reaches zero. The possibility of writing conditional statements and loops, and obtaining information like the magnetization state allows to construct arbitrarily complex simulation flows.

7. Conclusions and Outlook

MuMax is the first GPU-based micromagnetic solver that is publicly available as open-source software. Due to the large number of considered interaction terms and the versatile geometrical options (e.g. periodic boundary conditions) the software covers many of the classical micromagnetic research topics. The code is extensively validated by considering several standard problems and is shown to be reliable. The time gains are extremely large compared to CPU simulations: for large simulations a speedup with a factor 100 is easily obtained. These enormous speedups will open up new opportunities in micromagnetic modelling and boost fundamental magnetic research.

In the future, MuMax will be extended towards a yet more multipurpose software package incorporating other interactions in the Landau-Lifshitz equation: other expressions for the exchange contribution, exchange bias, magnetoelastic coupling, etc. Boundary correction methods to help in approximation non-square geometries better are currently also being considered. Furthermore, the description of more complex, non uniform microstructures will be made possible as e.g. nanocrystalline materials. Modules for hysteresis research and magnetic domain studies will be developed following the presented 2D and 3D approach as well as the so-called 2.5D approach (infinitely thick geometries).

Furthermore, the efficient simulation of yet larger problems is planned by introducing multiple GPUs –in one or more machines– for one single simulation. This way, the up to now limited memory available on the GPU hardware can be circumvented and the computational power will be further increased. Apart from requiring efficient communication between the different GPUs, this should be possible without drastic changes to our code as most of the implementation is already hardware-independent.

Acknowledgements

Financial support from the Flanders Research Foundation (FWO) is gratefully acknowledged. We cordially thank Bartel Van Waeyenberge, Luc Dupré, and Daniël De Zutter for supporting this research. Furthermore we would also like to thank André Drews, Claas Abert, Gunnar Selke and Theo Gerhardt from Hamburg University for the fruitful discussions and feedback.

References

References

- [1] MJ Donahue and DG Porter. OOMMF user's guide, version 1.0. interagency report NISTIR 6376, national institute of standards and technology, gaithersburg, MD, 1999.
- [2] W. Scholz, J. Fidler, T. Schrefl, D. Suess, R. Dittrich, H. Forster, and V. Tsiantos. Scalable parallel micromagnetic solvers for magnetic nanostructures. *Comput. Mater. Sci.*, 28:366–383, 2003.
- [3] T. Fischbacher, M. Franchin, G. Bordignon, and H. Fangohr. A systematic approach to multiphysics extensions of finite-element-based micromagnetic simulations: Nmag. *Ieee Transactions On Magnetism*, 43(6):2896–2898, 2007.
- [4] A. Kakay, E. Westphal, and R. Hertel. Speedup of fem micromagnetic simulations with graphical processing units. *Ieee Transactions On Magnetism*, 46(6):2303–2306, 2010.
- [5] S. J. Li, B. Livshitz, and V. Lomakin. Graphics processing unit accelerated o(n) micromagnetic solver. *Ieee Transactions On Magnetism*, 46(6):2373–2375, 2010.
- [6] G. Selke, A. Drews, and D.P.F. Möller. Highly efficient micromagnetic simulations using graphics processing units. *submitted to IEEE Trans. Magn.*, 2011.
- [7] B. Van de Wiele, F. Olyslager, and L. Dupre. Application of the fast multipole method for the evaluation of magnetostatic fields in micromagnetic computations. *Journal of Computational Physics*, 227(23):9913–9932, 2008.
- [8] M. J. Donahue and D. G. Porter. Exchange energy formulations for 3d micromagnetics. *Physica B-condensed Matter*, 343(1-4):177–183, 2004.
- [9] W. F. Brown Jr. *Micromagnetics*. Interscience Publishers, New York, NY, 1963.
- [10] S. Bohlens, B. Kruger, A. Drews, M. Bolte, G. Meier, and D. Pfannkuche. Current controlled random-access memory based on magnetic vortex handedness. *Applied Physics Letters*, 93(14):142508, October 2008.
- [11] S. S. P. Parkin, M. Hayashi, and L. Thomas. Magnetic domain-wall race-track memory. *Science*, 320(5873):190–194, April 2008.
- [12] L. Berger. Emission of spin waves by a magnetic multilayer traversed by a current. *Physical Review B*, 54(13):9353–9358, October 1996.
- [13] S. Zhang and Z. Li. Roles of nonequilibrium conduction electrons on the magnetization dynamics of ferromagnets. *Physical Review Letters*, 93(12):127204, September 2004.
- [14] B. Van de Wiele, F. Olyslager, and L. Dupre. Fast semianalytical time integration schemes for the landau-lifshitz equation. *Ieee Transactions On Magnetism*, 43(6):2917–2919, 2007.
- [15] M. Frigo and S. G. Johnson. The design and implementation of FFTW3. *Proceedings of the Ieee*, 93(2):216–231, 2005.
- [16] muMAG Micromagnetic Modeling Activity Group <http://www.ctcms.nist.gov/rdm/mumag.org.html>.
- [17] R. D. McMichael, M. J. Donahue, D. G. Porter, and J. Eicke. Comparison of magnetostatic field calculation methods on two-dimensional square grids as applied to a micromagnetic standard problem. *Journal of Applied Physics*, 85(8):5816–5818, 1999.
- [18] M. J. Donahue, D. G. Porter, R. D. McMichael, and J. Eicke. Behavior of mu mag standard problem no. 2 in the small particle limit. *Journal of Applied Physics*, 87(9):5520–5522, 2000.
- [19] M. Najafi, B. Kruger, S. Bohlens, M. Franchin, H. Fangohr, A. Vanhaverbeke, R. Allenspach, M. Bolte, U. Merkt, D. Pfannkuche, D. P. F. Moller,

and G. Meier. Proposal for a standard problem for micromagnetic simulations including spin-transfer torque. *Journal of Applied Physics*, 105(11), 2009.